

Sparse image reconstruction on the sphere: implications of a new sampling theorem

Jason D. McEwen, *Member, IEEE*, Gilles Puy, Jean-Philippe Thiran, *Senior Member, IEEE*,
Pierre Vandergheynst, Dimitri Van De Ville, *Senior Member, IEEE*, and Yves Wiaux, *Member, IEEE*

Abstract—We study the impact of sampling theorems on the fidelity of sparse image reconstruction on the sphere. We discuss how a reduction in the number of samples required to represent all information content of a band-limited signal acts to improve the fidelity of sparse image reconstruction, through both the dimensionality and sparsity of signals. To demonstrate this result we consider a simple inpainting problem on the sphere and consider images sparse in the magnitude of their gradient. We develop a framework for total variation (TV) inpainting on the sphere, including fast methods to render the inpainting problem computationally feasible at high-resolution. Recently a new sampling theorem on the sphere was developed, reducing the required number of samples by a factor of two for equiangular sampling schemes. Through numerical simulations we verify the enhanced fidelity of sparse image reconstruction due to the more efficient sampling of the sphere provided by the new sampling theorem.

Index Terms—Spheres, harmonic analysis, sampling methods, compressive sensing.

I. INTRODUCTION

IMAGES are observed on a spherical manifold in many fields, from astrophysics (e.g. [1]) and biomedical imaging (e.g. [2]), to computer graphics (e.g. [3]) and beyond. In many of these settings inverse problems arise, where one seeks to recover an unknown image from linear measurements, which may be noisy, incomplete or acquired through a convolution process, for example. Such inverse problems are typically

solved by assuming some regularising prior on the unknown image to be recovered.

Sparsity priors have received a lot of attention recently, since (a) they have been shown to be an effective and versatile approach for representing many real-world signals and (b) a sound theoretical foundation is provided by the emerging and rapidly evolving theory of compressive sensing [4]–[6]. On the sphere, compressive sensing has been considered for signals sparse in the spherical harmonic domain [7], however a general theoretical framework does not yet exist for signals sparse in spatially localised representations. Nevertheless, sparse image reconstruction on the sphere in alternative representations, such as a set of overcomplete dictionaries, may still be considered; indeed, such an approach has been shown to be very effective [8].

Although compressive sensing goes beyond Nyquist sampling, the Nyquist limit nevertheless defines the benchmark from which compressive sensing improvements are relative. Compressive sensing results are thus tightly coupled to the underlying sampling theorem on the manifold of interest. On the sphere, unlike Euclidean space, the number of samples required in the harmonic and spatial domains differ, with different sampling theorems on the sphere requiring a different number of samples in the spatial domain. Consequently, the sampling theorem adopted influences the performance of sparse signal reconstruction on the sphere. Studying the impact of sampling theorems on the sphere on the performance of sparse signal reconstruction is the focus of the current article.

When considering signal priors that incorporate spatially localised information (for example directly in real space, in the magnitude of the gradient of signals, or through a wavelet basis or overcomplete dictionary), the sampling theorem that is adopted becomes increasingly important. Recently, a new sampling theorem on the sphere was developed by two of the authors of the current article for equiangular sampling schemes [9], reducing Nyquist sampling on the sphere by a factor of two compared to the canonical approach [10], [11]. The reduction in the number of samples required to represent a band-limited signal on the sphere has important implications for sparse image reconstruction.

To gain some intuition regarding these implications, we appeal to standard compressive sensing results in Euclidean space, where the ratio of the number of measurements M required to reconstruct a sparse image, to its dimensionality N , goes as $M/N \propto K$ [5], [12], where K is the sparsity measure of the image (*i.e.* the number of non-zero coefficients

Manuscript received August 7, 2012; revised December 20, 2012; accepted February 15, 2013. Date of publication February 26 2013; date of current version April 12, 2013. We gratefully acknowledge use of some of the data and software made available on Frederik Simons' webpage: <http://www.frederik.net>. The work of J. D. McEwen is supported by the Swiss National Science Foundation (SNSF) under grant 200021-130359 and by a Newton International Fellowship from the Royal Society and the British Academy. The work of Y. Wiaux is supported in part by the Center for Biomedical Imaging (CIBM) of the Geneva and Lausanne Universities, Ecole Polytechnique Fédérale de Lausanne (EPFL), and the Leenaards and Louis-Jeantet foundations, and in part by the SNSF under grant PP00P2-123438.

Copyright (c) 2013 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

J. D. McEwen is with the Department of Physics and Astronomy, University College London, London WC1E 6BT, UK. G. Puy, J.-Ph. Thiran, P. Vandergheynst and Y. Wiaux are with the Signal Processing Laboratories, Institute of Electrical Engineering, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland. D. Van De Ville is with the Institute of Bioengineering, EPFL, CH-1015 Lausanne, Switzerland, and with the Department of Radiology and Medical Informatics, University of Geneva (UniGE), CH-1211 Geneva, Switzerland. G. Puy is also with the Institute of the Physics of Biological Systems, EPFL, CH-1015 Lausanne, Switzerland. Y. Wiaux is also with the Institute of Bioengineering, EPFL, CH-1015 Lausanne, Switzerland, and with the Department of Radiology and Medical Informatics, University of UniGE, CH-1211 Geneva, Switzerland.

E-mail: jason.mcewen@ucl.ac.uk (J. D. McEwen)

in some sparse representation).¹ If one is not concerned with the number of measurements required to achieve a given reconstruction fidelity but rather with the best fidelity for a given number of measurements, then this suggests reconstruction fidelity improves with decreasing dimensionality of the signal N and with decreasing sparsity K .

Both of these quantities, dimensionality and sparsity, are related to the number of samples required to capture all information content of the underlying signal, as prescribed by the adopted sampling theorem. Spatial dimensionality is given identically by the number of samples of the sampling theorem. For any sparse representation of an image that captures spatially localised information, the sparsity of the signal is also directly related to spatial sampling. For example, in a wavelet representation, wavelets are located on each sample point. A less dense dictionary of wavelet atoms required to span the space will lead to a more sparse representation of images when the sparsity is computed in an analysis approach, *i.e.* as the number of non-zero projections of the signal onto the wavelet atoms. We concentrate on such analysis priors here, as suggested by the recent evolution of compressive sensing with redundant dictionaries [12]. This argument can be extended to sparsity in the gradient and, in fact, all sparsity measures that capture spatially localised signal content. Consequently, for images sparse in a spatially localised representation, the ability to represent a band-limited signal on the sphere with fewer samples while still capturing all of its information content will improve the fidelity of sparse image reconstruction by enhancing both the dimensionality and sparsity of signals.

In this article we study the implications of a new sampling theorem [9] for sparse image reconstruction on the sphere. We verify the hypothesis that a more efficient sampling of the sphere, as afforded by the new sampling theorem [9], enhances the fidelity of sparse image reconstruction through both the dimensionality and sparsity of signals. To demonstrate this result we consider a simple inpainting problem, where we recover an image on the sphere from incomplete spatial measurements. We consider images sparse in the magnitude of their gradient, as an illustration of the general setting, and develop a framework for total variation (TV) inpainting on the sphere. Solving these problems is computationally challenging; hence we develop fast methods for this purpose. Our framework is general and is trivially extended to other sparsity priors that incorporate spatially localised information.

The remainder of the article is structured as follows. In Section II we concisely review the harmonic structure of the sphere and corresponding sampling theorems. We develop a framework for TV inpainting on the sphere in Section III. In Section IV we describe algorithms for solving the optimisation problems on the sphere that arise in our TV inpainting framework. Numerical simulations are performed in Section V,

showing the enhanced fidelity of sparse image reconstruction provided by a more efficient sampling of the sphere. Concluding remarks are made in Section VI.

II. SAMPLING ON THE SPHERE

A sampling theorem on the sphere states that all information in a (continuous) band-limited signal is captured in a finite number of samples in the spatial domain. Since a (continuous) band-limited signal on the sphere may be represented by a finite harmonic expansion, a sampling theorem on the sphere is equivalent to an exact prescription for computing a spherical harmonic transform from a finite set of spatial samples. In this section we review the harmonic structure of the sphere, before discussing sampling theorems on the sphere.

A. Harmonic structure of the sphere

We consider the space of square integrable functions on the sphere $L^2(S^2)$, with the inner product of $x, y \in L^2(S^2)$ defined by

$$\langle x, y \rangle \equiv \int_{S^2} d\Omega(\theta, \varphi) x(\theta, \varphi) y^*(\theta, \varphi),$$

where $d\Omega(\theta, \varphi) = \sin \theta d\theta d\varphi$ is the usual invariant measure on the sphere and (θ, φ) denote spherical coordinates with colatitude $\theta \in [0, \pi]$ and longitude $\varphi \in [0, 2\pi)$. Complex conjugation is denoted by the superscript $*$. The canonical basis for the space of square integrable functions on the sphere is given by the spherical harmonics $Y_{\ell m} \in L^2(S^2)$, with natural $\ell \in \mathbb{N}$, integer $m \in \mathbb{Z}$ and $|m| \leq \ell$. Due to the orthogonality and completeness of the spherical harmonics, any square integrable function on the sphere $x \in L^2(S^2)$ may be represented by its spherical harmonic expansion

$$x(\theta, \varphi) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} \hat{x}_{\ell m} Y_{\ell m}(\theta, \varphi), \quad (1)$$

where the spherical harmonic coefficients are given by the usual projection onto each basis function:

$$\hat{x}_{\ell m} = \langle x, Y_{\ell m} \rangle = \int_{S^2} d\Omega(\theta, \varphi) x(\theta, \varphi) Y_{\ell m}^*(\theta, \varphi).$$

Throughout, we consider signals on the sphere band-limited at L , that is signals such that $\hat{x}_{\ell m} = 0, \forall \ell \geq L$, in which case the summation over ℓ in (1) may be truncated to the first L terms. Finally, note that the harmonic coefficients of a real function on the sphere satisfy the conjugate symmetry relation $\hat{x}_{\ell m}^* = (-1)^m \hat{x}_{\ell, -m}$, which follows directly from the conjugate symmetry of the spherical harmonics.

B. Sampling theorems on the sphere

Sampling theorems on the sphere describe how to sample a band-limited signal x so that all information is contained in a finite number of samples N . Moreover, a sampling theorem on the sphere effectively encodes an exact quadrature rule for the integration of band-limited functions [9], [10]. We denote the concatenated vector of N spatial measurements by $\mathbf{x} \in \mathbb{C}^N$ and the concatenated vector of L^2 harmonic coefficients by $\hat{\mathbf{x}} \in \mathbb{C}^{L^2}$. The number of spatial and harmonic elements, N

¹Typically the mutual coherence of the measurement and sparsifying operators also plays a role [5]. However, in Euclidean space, as on the sphere, discrete inner products can be related to the (unique) continuous inner product (via a sampling theorem). Consequently, the measure of coherence is invariant to the choice of sampling theorem (the coherence is defined through the continuous inner product, which is the same for all sampling theorems). For the purpose of comparing sampling theorems on the sphere, we can thus safely neglect the impact of coherence.

and L^2 respectively, may differ (and in fact do differ for all known sampling theorems on the sphere).

Before discussing different sampling theorems on the sphere, we define a generic notation to describe the harmonic transform corresponding to a given sampling theorem. A sampling theorem describes how to compute the spherical harmonic transform of a signal exactly. Since the spherical harmonic transform and inverse are linear, we represent the forward and inverse transform by the matrix operators $\Gamma \in \mathbb{C}^{L^2 \times N}$ and $\Lambda \in \mathbb{C}^{N \times L^2}$ respectively. The spherical harmonic coefficients of a sampled signal (*i.e.* image) on the sphere \mathbf{x} are given by the forward transform

$$\hat{\mathbf{x}} = \Gamma \mathbf{x},$$

while the original signal is recovered from its harmonic coefficients by the inverse transform

$$\mathbf{x} = \Lambda \hat{\mathbf{x}}.$$

Different sampling theorems then differ in the definition of Λ , Γ and the number of spatial samples N . By definition, all sampling theorems give exact spherical harmonic transforms, implying $\Gamma\Lambda = \mathbb{1}_{L^2}$, where $\mathbb{1}_k$ is the $k \times k$ identity matrix. However, for all sampling theorems on the sphere the number of samples required in the spatial domain exceeds the number of coefficients in the harmonic domain (*i.e.* $N > L^2$), hence $\Lambda\Gamma \neq \mathbb{1}_N$. Consequently, for the N sample positions of a sampling theorem, an arbitrary set of sample values does not necessarily define a band-limited signal (contrast this to the discrete Euclidean setting where a finite set of samples uniquely defines a band-limited signal). Note also that the adjoint inverse (forward) spherical harmonic transform differs to the forward (inverse) spherical harmonic transform in the discrete setting.

For an equiangular sampling of the sphere, the Driscoll & Healy (DH) [10] sampling theorem has become the standard, requiring $N_{\text{DH}} = 2L(2L - 1) \sim 4L^2$ samples on the sphere to represent exactly a signal band-limited in its spherical harmonic decomposition at L . Recently, a new sampling theorem for equiangular sampling schemes has been developed by McEwen & Wiaux (MW) [9], requiring only $N_{\text{MW}} = (L - 1)(2L - 1) + 1 \sim 2L^2$ samples to represent a band-limited signal exactly. No sampling theorem on the sphere reaches the optimal number of samples suggested by the L^2 dimension of a band-limited signal in harmonic space (although the MW sampling theorem comes closest to this bound). The MW sampling theorem therefore achieves a more efficient sampling of the sphere, with a reduction by a factor of approximately two in the number of samples required to represent a band-limited signal on the sphere.²

²Gauss-Legendre (GL) quadrature can also be used to construct an efficient sampling theorem on the sphere, with $N_{\text{GL}} = L(2L - 1) \sim 2L^2$ samples (see *e.g.* [9]). The MW sampling theorem nevertheless remains more efficient, especially at low band-limits. Furthermore, it is not so straightforward to define the TV norm on the GL grid since it is not equiangular. Finally, algorithms implementing the GL sampling theorem have been shown to be limited to lower band-limits and less accurate than the algorithms implementing the MW sampling theorem [9]. Thus, we focus on equiangular sampling theorems only in this article.

Fast algorithms have been developed to compute forward and inverse spherical harmonic transforms rapidly for both the DH [10], [11] and MW [9] sampling theorems. These fast algorithms are implemented, respectively, in the publicly available `SpharmonicKit`³ package and the `Spin Spherical Harmonic Transform (SSHT)`⁴ package and are essential to facilitate the application of these sampling theorems at high band-limits.

III. SPARSE IMAGE RECONSTRUCTION ON THE SPHERE

A more efficient sampling of a band-limited signal on the sphere, as afforded by the MW sampling theorem, improves the quality of sparse image reconstruction for images that are sparse in a spatially localised measure. To demonstrate this result we consider a simple inpainting problem on the sphere and consider images sparse in the magnitude of their gradient. We develop a framework for total variation (TV) inpainting on the sphere, which relies on a sampling theorem and its associated quadrature rule to define a discrete TV norm on the sphere. Firstly, we define the discrete TV norm on the sphere, before secondly defining finite difference gradient operators on the sphere. Thirdly, we discuss the TV inpainting problem.

A. TV norm on the sphere

We define the discrete TV norm on the sphere by

$$\|\mathbf{x}\|_{\text{TV}} \equiv \sum_{t=0}^{N_\theta-1} \sum_{p=0}^{N_\varphi-1} q(\theta_t) |\nabla \mathbf{x}|, \quad (2)$$

where t and p index the equiangular samples in θ and φ respectively, with the number of samples associated with a given sampling theorem denoted in each dimension by N_θ and N_φ respectively. The discrete magnitude of the gradient is defined by

$$|\nabla \mathbf{x}| \equiv \sqrt{(\delta_\theta \mathbf{x})^2 + \frac{1}{\sin^2 \theta} (\delta_\varphi \mathbf{x})^2}, \quad (3)$$

to approximate the continuous magnitude of the gradient

$$|\nabla x| \equiv \sqrt{\left(\frac{\partial x}{\partial \theta}\right)^2 + \frac{1}{\sin^2 \theta} \left(\frac{\partial x}{\partial \varphi}\right)^2}$$

by finite differences. The finite difference operators δ_θ and δ_φ are defined explicitly in the following subsection. The contribution to the TV norm from the magnitude of the gradient for each pixel value is weighted by the quadrature weights $q(\theta_t)$ of the sampling theorem adopted in order to approximate

³<http://www.cs.dartmouth.edu/~geelong/sphere/>

⁴<http://www.jasonmcewen.org/>

continuous integration.⁵ The inclusion of the weights $q(\theta_t)$ also regularises the $\sin \theta$ term that arises from the definition of the gradient on the sphere, eliminating numerical instabilities that this would otherwise cause.

B. Gradient operators on the sphere

The finite difference operators δ_θ and δ_φ defined on the sphere appear in the definition of the discrete magnitude of the gradient given by (3), and thus are required to compute the discrete TV norm on the sphere. Furthermore, as we shall see, to solve the TV inpainting problems outlined in the following subsection, the adjoints of these operators are also required. We define these operators and adjoints explicitly here.

The operator δ_θ is defined sample-wise by

$$\mathbf{u}_{t,p} \equiv (\delta_\theta \mathbf{x})_{t,p} \equiv \begin{cases} \mathbf{x}_{t+1,p} - \mathbf{x}_{t,p}, & t = 0, 1, \dots, N_\theta - 2 \text{ and } \forall p \\ 0, & t = N_\theta - 1 \text{ and } \forall p \end{cases},$$

with adjoint

$$(\delta_\theta^\dagger \mathbf{u})_{t,p} = \begin{cases} -\mathbf{u}_{t,p}, & t = 0 \text{ and } \forall p \\ \mathbf{u}_{t-1,p} - \mathbf{u}_{t,p}, & t = 1, \dots, N_\theta - 2 \text{ and } \forall p \\ \mathbf{u}_{t-1,p}, & t = N_\theta - 1 \text{ and } \forall p \end{cases}.$$

Note that this definition is identical to the typical definition of the corresponding operator on the plane [13]. The operator δ_φ is defined sample-wise by

$$\mathbf{v}_{t,p} \equiv (\delta_\varphi \mathbf{x})_{t,p} \equiv \begin{cases} \mathbf{x}_{t,p+1} - \mathbf{x}_{t,p}, & p = 0, 1, \dots, N_\varphi - 2 \text{ and } \forall t \\ \mathbf{x}_{t,0} - \mathbf{x}_{t,p}, & p = N_\varphi - 1 \text{ and } \forall t \end{cases},$$

with adjoint

$$(\delta_\varphi^\dagger \mathbf{v})_{t,p} = \begin{cases} \mathbf{v}_{t,N_\varphi-1} - \mathbf{v}_{t,p}, & p = 0 \text{ and } \forall t \\ \mathbf{v}_{t,p-1} - \mathbf{v}_{t,p}, & p = 1, \dots, N_\varphi - 1 \text{ and } \forall t \end{cases}.$$

Since the sphere is periodic in φ , we define the corresponding finite difference operator to also be periodic. The finite difference operator and adjoint in φ therefore differ to the typical definition on the plane [13].

The TV norm on the sphere may then be seen as the sum of the magnitude of the weighted gradient

$$\|\mathbf{x}\|_{\text{TV}} = \sum_{t=0}^{N_\theta-1} \sum_{p=0}^{N_\varphi-1} |(\tilde{\nabla} \mathbf{x})_{t,p}|,$$

where

$$|(\tilde{\nabla} \mathbf{x})_{t,p}| = (\tilde{\mathbf{u}}_{t,p}^2 + \tilde{\mathbf{v}}_{t,p}^2)^{1/2},$$

⁵If the band-limiting operator $\Upsilon \equiv \Lambda \Gamma \in \mathbb{C}^{N \times N}$ were applied to $|\nabla \mathbf{x}|$ in (2), then the finite summation of (2) would give an exact quadrature for the integral of the continuous function underlying the associated samples of the band-limited $|\nabla \mathbf{x}|$. However, introducing the operator Υ in (2) would make solving the optimisation problems defined subsequently problematic and would also prohibit passing the quadrature weights inside the gradient to eliminate numerical instabilities due to the $\sin \theta$ term. Consequently, we adopt the definition of the discrete TV norm on the sphere given by (2). In any case, numerical experiments have shown that $\|\mathbf{x}\|_{\text{TV}}$ is identical to $\sum_{t,p} q(\theta_t) \Upsilon |\nabla \mathbf{x}|$ to machine precision for the particular test images considered in Section V-A. Thus, the discrete TV norm defined by (2) can be thought of as an accurate proxy for $\int_{S^2} d\Omega \Upsilon |\nabla \mathbf{x}|$.

for

$$\begin{pmatrix} \tilde{\mathbf{u}} \\ \tilde{\mathbf{v}} \end{pmatrix} \equiv \tilde{\nabla} \mathbf{x}.$$

The weighted gradient operator is defined by

$$\tilde{\nabla} \equiv \begin{pmatrix} \tilde{\delta}_\theta \\ \tilde{\delta}_\varphi \end{pmatrix},$$

where the weighted finite difference operators are defined by

$$(\tilde{\delta}_\theta)_{t,p} \equiv q(\theta_t) (\delta_\theta)_{t,p}$$

and

$$(\tilde{\delta}_\varphi)_{t,p} \equiv \frac{q(\theta_t)}{\sin \theta_t} (\delta_\varphi)_{t,p}.$$

Notice how the inclusion of the weights $q(\theta_t)$ regularises the $\sin \theta$ term that arises from the definition of the gradient on the sphere, eliminating numerical instabilities that this would otherwise cause. If $\theta_t = \pi$, corresponding to the South pole of the sphere, then $(\delta_\varphi \mathbf{x})_{t,p} = 0$ and thus we define $(\tilde{\delta}_\varphi \mathbf{x})_{t,p} = 0$ to avoid dividing by $\sin \theta_t = 0$. Note that the MW sampling theorem includes a sample on the South pole, while the DH sampling theorem does not (neither sampling theorem includes a sample on the North pole). The adjoint weighted gradient operator is then applied as

$$\mathbf{x}' = \tilde{\nabla}^\dagger \begin{pmatrix} \tilde{\mathbf{u}} \\ \tilde{\mathbf{v}} \end{pmatrix} = \tilde{\delta}_\theta^\dagger \tilde{\mathbf{u}} + \tilde{\delta}_\varphi^\dagger \tilde{\mathbf{v}},$$

where the adjoint operators $\tilde{\delta}_\theta^\dagger$ and $\tilde{\delta}_\varphi^\dagger$ follow trivially from δ_θ^\dagger and δ_φ^\dagger .

C. TV inpainting on the sphere

We consider the measurement equation

$$\mathbf{y} = \Phi \mathbf{x} + \mathbf{n},$$

where M noisy real measurements $\mathbf{y} \in \mathbb{R}^M$ of the underlying real image on the sphere $\mathbf{x} \in \mathbb{R}^N$ are made. The matrix implementing the measurement operator $\Phi \in \mathbb{R}^{M \times N}$ represents a uniformly random masking of the image, with one non-zero, unit value on each row specifying the location of the measured datum. The noise $\mathbf{n} \in \mathbb{R}^M$ is assumed to be independent and identically distributed (iid) Gaussian noise, with zero mean and variance σ_n^2 . We assume that the image \mathbf{x} is sparse in the norm of its gradient and thus attempt to recover \mathbf{x} from measurements \mathbf{y} by solving the following TV inpainting problem directly on the sphere:

$$\mathbf{x}^* = \arg \min_{\mathbf{x}} \|\mathbf{x}\|_{\text{TV}} \text{ such that } \|\mathbf{y} - \Phi \mathbf{x}\|_2 \leq \epsilon. \quad (4)$$

The square of the residual noise follows a scaled χ^2 distribution with M degrees of freedom, *i.e.* $\|\mathbf{y} - \Phi \mathbf{x}^*\|_2^2 \sim \sigma_n^2 \chi^2(M)$. Consequently, we choose ϵ^2 to correspond to the $(100\alpha)\text{th}$ percentile of this distribution, giving a probability α that pure noise produces a residual noise equal to or smaller than the observed residual. Note that the data constraint in (4) is given by the usual ℓ_2 -norm, which is appropriate for Gaussian noise on a discrete set of measurements. Although we consider band-limited signals, we have not imposed this constraint when

solving (4). Consequently, \mathbf{x}^* will not necessarily be band-limited at L and we impose this constraint on the solution by performing a forward and inverse spherical harmonic transform: $\mathbf{x}_L^* = \Upsilon \mathbf{x}^*$, where the band-limiting operator is defined by $\Upsilon \equiv \Lambda \Gamma \in \mathbb{C}^{N \times N}$.

As discussed already, for images sparse in a measure that captures spatially localised information, such as the TV norm, a more efficient sampling of the signal enhances sparsity. Furthermore, when recovering signals in the spatial domain directly, the dimensionality of the signal is also enhanced by a more efficient sampling. These two effects both act to improve the fidelity of sparse image reconstruction. Thus, the more efficient sampling of the MW sampling theorem when compared to the DH sampling theorem will improve the fidelity of sparse image reconstruction when solving the TV inpainting problem given by (4). We verify these claims with numerical experiments in Section V.

No sampling theorem on the sphere reaches the optimal number of samples in the spatial domain suggested by the L^2 dimensionality of the signal in the harmonic domain. We may therefore optimise the dimensionality of the signal that we attempt to recover by recovering its harmonic coefficients $\hat{\mathbf{x}}$ directly. We do so by solving the following TV inpainting problem in harmonic space:

$$\hat{\mathbf{x}}^* = \arg \min_{\hat{\mathbf{x}}'} \|\Lambda' \hat{\mathbf{x}}'\|_{\text{TV}} \text{ such that } \|\mathbf{y} - \Phi \Lambda' \hat{\mathbf{x}}'\|_2 \leq \epsilon. \quad (5)$$

We impose reality of the recovered signal by explicitly imposing conjugate symmetry in harmonic space through the conjugate symmetry extension operator $\Pi \in \mathbb{C}^{L^2 \times L(L+1)/2}$, where $\Lambda' = \Lambda \Pi$. The full set of harmonic coefficients of \mathbf{x} are given by $\hat{\mathbf{x}} = \Pi \hat{\mathbf{x}}'$, where $\hat{\mathbf{x}}' \in \mathbb{C}^{L(L+1)/2}$ are the harmonic coefficients for the spherical harmonic azimuthal index m non-negative only. The image on the sphere is then recovered from its harmonic coefficients by $\mathbf{x}^* = \Lambda' \hat{\mathbf{x}}^*$. By solving the TV inpainting problem directly in harmonic space, we naturally recover a signal band-limited at L .

When solving the TV inpainting problem (5) directly in harmonic space, the dimensionality of the recovered signal is optimal and identical for both sampling theorems. However, the sparsity of the signal with respect to the TV norm remains enhanced for the MW sampling theorem when compared to the DH sampling theorem. Consequently, the MW sampling theorem will improve the fidelity of sparse image reconstruction when solving the TV inpainting problem given by (5), although through sparsity only and not also dimensionality. We verify these claims with numerical experiments in Section V.

Note that if a band-limit constraint were explicitly imposed in problem (4), then the two problems would be equivalent, however, this would involve applying the band-limiting operator $\Upsilon = \Lambda \Gamma$, complicating the problem and increasing the computational cost of finding a solution, while providing no improvement over (5). In the current formulation of these two optimisation problems, problem (4) has the advantage of simplicity, while problem (5) is the simplest formulation that optimises dimensionality.

IV. ALGORITHMS

We solve the TV inpainting problems on the sphere given by (4) and (5) using iterative convex optimisation methods. Solving the TV inpainting problem in harmonic space poses two challenges as we go to high band-limits (*i.e.* high-resolution). Firstly, we require as an input to the convex optimisation algorithm an upper bound on the inverse transform operator norm, which is challenging to compute at high-resolution. We describe a method to compute the operator norm at high-resolution, which, crucially, does not require an explicit computation of Λ . Secondly, the inverse spherical harmonic transform Λ and its adjoint operator Λ^\dagger must be applied repeatedly in the iterative algorithm. Fast algorithms are essential to perform forward and inverse spherical harmonic transforms at high-resolution and have been developed for both the DH [10], [11] and MW [9] sampling theorems. To solve the inpainting problem at high-resolution we also require a fast adjoint inverse transform. We thus develop fast algorithms to perform the adjoint forward and adjoint inverse spherical harmonic transforms corresponding to the MW sampling theorem. Since we predict the MW sampling theorem to be superior to the DH sampling theorem for sparse image reconstruction on the sphere (a prediction that is validated by numerical experiments performed at low band-limits (*i.e.* low-resolution) in Section V), we develop fast adjoint algorithms for the MW sampling theorem only. These methods then render the computation of solutions to the TV inpainting problems feasible at high-resolution for the MW sampling theorem.

A. Convex optimisation

We apply the Douglas-Rachford proximal splitting algorithm [14] to solve the convex optimisation problems (4) and (5).⁶ We describe only how to solve problem (5) as problem (4) can be solved in the same way (by replacing Λ' with the identity matrix $\mathbb{1}_N$ and by replacing $\hat{\mathbf{x}}'$ with \mathbf{x}).

The Douglas-Rachford algorithm [14] is based on a splitting approach that requires the computation of two proximity operators [15]. In our case, one proximity operator is based on the TV norm $\|\Lambda' \cdot\|_{\text{TV}}$ and the other on the data constraint $\|\mathbf{y} - \Phi \Lambda' \cdot\|_2 \leq \epsilon$.

In the case of an image on the plane, the proximity operator based on the TV norm may be computed using, for example, the method described in [13] or in [16]. For an image on the sphere, the same methods can be used after introducing the following modifications. In [16] the algorithm to compute the proximity operator of the TV norm is described in terms of a linear operator \mathcal{L} , its adjoint \mathcal{L}^\dagger , and two projections onto a set \mathcal{P} and a set \mathcal{C} . In our case, the linear operator \mathcal{L} and its adjoint \mathcal{L}^\dagger may be redefined as

$$\mathcal{L} : \begin{pmatrix} \tilde{\mathbf{u}} \\ \tilde{\mathbf{v}} \end{pmatrix} \mapsto -\Lambda'^\dagger \tilde{\nabla}^\dagger \begin{pmatrix} \tilde{\mathbf{u}} \\ \tilde{\mathbf{v}} \end{pmatrix}$$

⁶We use Douglas-Rachford splitting since this does not require differentiability of the objective function and allows us to solve constrained optimisation problems where we adopt an indicator function to represent the measurement constraint.

and

$$\mathcal{L}^\dagger : \hat{\mathbf{x}}' \mapsto -\tilde{\nabla} \Lambda' \hat{\mathbf{x}}' = - \begin{pmatrix} \tilde{\delta}_\theta \Lambda' \hat{\mathbf{x}}' \\ \tilde{\delta}_\varphi \Lambda' \hat{\mathbf{x}}' \end{pmatrix},$$

where the set \mathcal{P} is the set of weighted gradient-pairs $(\tilde{\mathbf{u}}, \tilde{\mathbf{v}})$ such that $\tilde{\mathbf{u}}_{t,p}^2 + \tilde{\mathbf{v}}_{t,p}^2 \leq 1$ and \mathcal{C} is simply given by the space of the recovered vector $\hat{\mathbf{x}}$.

The second proximity operator, related to the data constraint $\|\mathbf{y} - \Phi \Lambda' \cdot\|_2 \leq \epsilon$, is computed using the method described in [17] directly.

B. Operator norm bound

The convex optimisation algorithm requires as input upper bounds for the norms of the operators that appear in the problem. The calculation of these norms is in most cases straightforward, however the calculation of the inverse spherical harmonic transform operator norm, defined by

$$\|\Lambda\|_2 \equiv \max_{\|\hat{\mathbf{x}}\|_2=1} \|\Lambda \hat{\mathbf{x}}\|_2,$$

can prove problematic. At low-resolution $\|\Lambda\|_2$ may be computed explicitly, however this is not feasible at high-resolution since even computing and storing Λ explicitly is challenging.

We develop a method here to estimate this norm for the MW sampling theorem without computing Λ explicitly. We seek a sampled function on the sphere $\mathbf{x} = \Lambda \hat{\mathbf{x}}$ that maximises $\|\mathbf{x}\|_2$, while satisfying the constraint $\|\hat{\mathbf{x}}\|_2 = 1$. By the Parseval relation and the sampling theorem on the sphere, this constraint may be rewritten:

$$\begin{aligned} \|\hat{\mathbf{x}}\|_2 = 1 & \quad \Rightarrow \quad \langle \mathbf{x}, \mathbf{x} \rangle = 1 \\ & \quad \text{Parseval} \\ & \quad \Rightarrow \quad \mathbf{x}_u^\dagger Q_u \mathbf{x}_u = 1, \\ & \quad \text{Sampling theorem} \end{aligned}$$

where $\mathbf{x}_u \in \mathbb{R}^{N_u}$ contains samples of \mathbf{x} , sampled at a resolution sufficient to represent \mathbf{x}^2 , *i.e.* corresponding to band-limit $2L-1$ (so that an exact quadrature may be used to evaluate $\langle \mathbf{x}, \mathbf{x} \rangle$ from a discrete set of samples), $Q_u \in \mathbb{R}^{N_u \times N_u}$ is the matrix with corresponding quadrature weights along its diagonal, and where $N_u \sim 2(2L-1)^2$. Since we know that the quadrature weights for the MW sampling theorem are closely approximated by $\sin \theta$ [9], the signal that maximises $\|\mathbf{x}\|_2$ while satisfying the constraint $\mathbf{x}_u^\dagger Q_u \mathbf{x}_u = 1$ has its energy centred as much as possible on the South pole since this is where the quadrature weights are smallest (recall that the MW sampling scheme does not contain a sample on the North pole). This signal is given by the band-limited Dirac delta function centred on the South pole (see *e.g.* [18] for the definition of the band-limited Dirac delta function on the sphere). The spherical harmonic coefficients of this band-limited Dirac delta function $\hat{\delta}^L \in L^2(S^2)$ are given by

$$\hat{\delta}_{\ell m}^L = \kappa (-1)^\ell \sqrt{\frac{2\ell+1}{4\pi}} \delta_{m0},$$

where κ is a normalisation factor chosen to ensure $\|\hat{\delta}^L\|_2 = 1$ and δ_{ij} is the Kronecker delta symbol. The norm of the inverse spherical harmonic transform operator may then be computed




Fig. 1. Explicit calculation of the inverse spherical harmonic transform operator norm $\|\Lambda\|_2$ and estimation by the method outlined in the text, at low-resolution. The solid red line shows the estimated norm for all band-limits L , while the solid blue circles show the values computed explicitly for $L \in \{2, 4, 8, 16, 32\}$. The estimated norm agrees with the actual norm very well.

by $\|\Lambda\|_2 \simeq \|\Lambda \hat{\delta}^L\|_2$, which, crucially, does not require an explicit computation of Λ , merely its application.

In Figure 1 we compute $\|\Lambda\|_2$ by the method outlined here and from Λ explicitly, for low-resolution. We find that the method to estimate the norm of the inverse spherical harmonic transform operator outlined here estimates the actual norm very well.

We also derived an upper bound for the norm of this operator for the MW sampling theorem. However, the bound we derived is not tight and we found empirically that the method outlined here to estimate the norm itself, rather than a bound, is very accurate and improved the performance of the optimisation algorithm considerably when compared to a non-tight bound. Although we do not prove so explicitly, we conjecture that the method outlined here gives the inverse transform operator norm exactly.

C. Fast adjoint spherical harmonic transforms

Standard convex optimisation methods require not only the application of the operators that appear in the optimisation problem but often also their adjoints. Moreover, these methods are typically iterative, necessitating repeated application of each operator and its adjoint. Thus, to solve optimisation problems that incorporate harmonic transform operators, like the harmonic space TV inpainting problem given by (5), fast algorithms to apply both the operator and its adjoint are required to render high-resolution problems computationally feasible.

Here we develop fast algorithms to perform adjoint forward and adjoint inverse spherical harmonic transforms for the MW sampling theorem. Although we only require the adjoint inverse transform in this article, for the sake of completeness we also derive a fast adjoint forward transform. Similarly, although we only consider scalar functions in this article, for the sake of completeness we derive fast adjoint algorithms

for the spin setting. A spin function on the sphere transforms as ${}_sf'(\theta, \varphi) = e^{-is\chi} {}_sf(\theta, \varphi)$ under a local rotation by $\chi \in [0, 2\pi)$, where the prime denotes the rotated function. It is important to note that the rotation considered here is *not* a global rotation on the sphere but rather a rotation by χ in the tangent plane at (θ, φ) (see *e.g.* [9] for further details). In the expressions for the fast algorithms derived below, the standard scalar case follows simply by setting $s = 0$. These fast adjoint algorithms are implemented in the publicly available SSHT⁷ package [9].

The fast adjoint inverse spherical harmonic transform for the MW sampling theorem follows by taking the adjoint of each stage of the fast inverse transform [9] and applying these in reverse order. The final stage of the fast inverse transform involves discarding out-of-domain samples and has adjoint

$${}_sf^\dagger(\theta_t, \varphi_p) = \begin{cases} {}_sf(\theta_t, \varphi_p), & t \in \{0, 1, \dots, L-1\} \\ 0, & t \in \{L, \dots, 2L-2\} \end{cases}.$$

The second stage of the fast adjoint inverse transform is given by

$${}_sF_{mm'}^\dagger = \sum_{t=0}^{2L-2} \sum_{p=0}^{2L-2} {}_sf^\dagger(\theta_t, \varphi_p) e^{-i(m'\theta_t + m\varphi_p)},$$

which may be computed rapidly using fast Fourier transforms (FFTs). The final stage of the fast adjoint inverse transform is given by

$$\begin{aligned} {}_sf_{\ell m}^\dagger &= (-1)^s i^{m+s} \sqrt{\frac{2\ell+1}{4\pi}} \\ &\times \sum_{m'=-\ell}^{\ell-1} \Delta_{m'm}^\ell \Delta_{m',-s}^\ell {}_sF_{mm'}^\dagger, \end{aligned}$$

where $\Delta_{mn}^\ell \equiv d_{mn}^\ell(\pi/2)$ are the Wigner d -functions evaluated for argument $\pi/2$ (see *e.g.* [19]). This final calculation dominates the overall asymptotic complexity of the fast adjoint inverse transform, resulting in an algorithm with complexity $\mathcal{O}(L^3)$.

The fast adjoint forward spherical harmonic transform for the MW sampling theorem follows by taking the adjoint of each stage of the fast forward transform [9] and applying these in reverse order. The first stage of the fast adjoint forward transform is given by

$$\begin{aligned} {}_sG_{mm'}^\dagger &= (-1)^s i^{-(m+s)} \\ &\times \sum_{\ell=0}^{L-1} \sqrt{\frac{2\ell+1}{4\pi}} \Delta_{m'm}^\ell \Delta_{m',-s}^\ell {}_sf_{\ell m}. \end{aligned}$$

The next stage is given by the (reflected) convolution

$${}_sF_{mm'}^\dagger = 2\pi \sum_{m''=-(L-1)}^{L-1} {}_sG_{mm'}^\dagger w(m' - m''),$$

which is self-adjoint, followed by the inverse Fourier transform in θ

$${}_s\tilde{F}_m^\dagger(\theta_t) = \frac{1}{2L-1} \sum_{m'=-\ell}^{\ell-1} {}_sF_{mm'}^\dagger e^{im'\theta_t},$$

which may be computed rapidly using FFTs. The next stage consists of the adjoint of the periodic extension of a function on the sphere performed in the forward transform and is given by

$${}_sF_m^\dagger(\theta_t) = \begin{cases} {}_s\tilde{F}_m^\dagger(\theta_t) \\ + (-1)^{m+s} {}_s\tilde{F}_m^\dagger(\theta_{2L-2-t}), & t \in \{0, 1, \dots, L-2\} \\ {}_s\tilde{F}_m^\dagger(\theta_t), & t = L-1 \end{cases}.$$

The final stage consists of the Fourier transform in φ

$${}_sf^\dagger(\theta_t, \varphi_p) = \frac{1}{2L-1} \sum_{m=-(L-1)}^{L-1} {}_sF_m^\dagger(\theta_t) e^{im\varphi_p},$$

which may be computed rapidly using FFTs. The first calculation dominates the overall asymptotic complexity of the fast adjoint forward transform, resulting in an algorithm with complexity $\mathcal{O}(L^3)$.

V. SIMULATIONS

We perform numerical experiments to examine the impact of a more efficient sampling of the sphere when solving the TV inpainting problems defined in Section III. Firstly, we perform a low-resolution comparison of reconstruction fidelity when adopting the DH and MW sampling theorems, where the predicted improvements in reconstruction fidelity provided by the MW sampling theorem are verified in practice. Secondly, we perform a single simulation to illustrate TV inpainting at high-resolution on a realistic test image.

A. Low-resolution comparison on band-limited images

A test image is constructed from Earth topography data. The original Earth topography data are taken from the Earth Gravitational Model (EGM2008) publicly released by the U.S. National Geospatial-Intelligence Agency (NGA) EGM Development Team.⁸ To create a band-limited test signal sparse in its gradient, the original data are thresholded at their midpoint to create a binary Earth map (scaled to contain zero and unit values), which is then smoothed by multiplication in harmonic space with the Gaussian $\hat{G}_{\ell m} = \exp(-\ell^2 \sigma_s)$, with $\sigma_s = 0.002$, to give a signal band-limited at $L = 32$. The resulting test image is displayed in Figure 2. Let us stress that this test image is constructed to satisfy the assumptions of our theoretical framework, *i.e.* the case of band-limited images that are sparse in their gradient. This is necessary to evaluate the theoretical predictions based on our framework. A realistic test image is considered in the following subsection.

Measurements of the test image are taken at uniformly random locations on the sphere, as described by the measurement operator Φ , in the presence of Gaussian iid noise with standard deviation $\sigma_n = 0.01$. Reconstructed images on the sphere are recovered by solving the inpainting problems in the spatial and harmonic domains, through (4) and (5) respectively, using both the DH and MW sampling theorems,

⁸These data were downloaded and extracted using the tools available from Frederik Simons' webpage: <http://www.frederik.net>.

⁷<http://www.jasonmcewen.org/>




Fig. 2. Test image of Earth topographic data constructed to be sparse in its gradient and band-limited at $L = 32$. This image constitutes the ground truth in our numerical experiments. Here and subsequently data on the sphere are displayed using the Mollweide projection, with zero values shown in black, unit values shown in yellow, and the colour of intermediate values interpolated between these extremes.

giving four reconstruction techniques. The bound ϵ is determined from $\alpha = 0.99$. We consider the measurement ratios $M/L^2 \in \{1/4, 1/2, 1, 3/2, N_{\text{MW}}/L^2 \sim 2\}$ (recall that L^2 is the dimensionality of the signal in harmonic space). The measurement ratio $M/L^2 = N_{\text{MW}}/L^2 \sim 2$ corresponds to complete coverage for the MW sampling theorem, *i.e.* Nyquist rate sampling on the MW grid.

Typical reconstructed images are shown in Figure 3 for the four reconstruction techniques. For each reconstruction technique and measurement ratio M/L^2 , we perform ten simulations for random measurement operators and noise. To quantify the error of reconstruction, we compute the signal-to-noise-ratio $\text{SNR} = 20 \log(\|\hat{\mathbf{x}}\|_2 / \|\hat{\mathbf{x}}^* - \hat{\mathbf{x}}\|_2)$ (defined in harmonic space to avoid differences due to the number of samples of each sampling theorem). Note that the standard ℓ_2 -norm is used in the definition of the SNR given the discrete nature of harmonic space on the sphere. Reconstruction performance, averaged over these ten simulations, is shown in Figure 4.

When solving the inpainting problem in the spatial domain through (4) we see a large improvement in reconstruction quality for the MW sampling theorem when compared to the DH sampling theorem. This is due to the enhancement in both dimensionality and sparsity afforded by the MW sampling theorem in this setting. When solving the inpainting problem in the harmonic domain through (5) we see a considerable improvement in reconstruction quality for each sampling theorem, since we optimise the dimensionality of the recovered signal by going to harmonic space. For harmonic reconstructions, the MW sampling theorem remains superior to the DH sampling theorem due to the enhancement in sparsity (but not dimensionality) that it affords in this setting. All of the predictions made in Section III are thus exhibited in the numerical experiments performed in this section. In all cases, the superior performance of the MW sampling theorem is clear.

B. High-resolution illustration on a realistic image

In this section we perform a single simulation to illustrate TV inpainting at high resolution. Furthermore, we also consider a more realistic test image. Since we develop fast adjoint algorithms for the MW sampling theorem only (due to its




Fig. 4. Reconstruction performance for the DH (green/diamonds) and MW (red/circles) sampling theorems, when solving the TV inpainting problem in the spatial (dot-dashed line) and harmonic domain (solid line). The MW sampling theorem provides enhancements in reconstruction quality when compared to the DH sampling theorem, due to dimensionality and sparsity improvements in spatial reconstructions, and due to sparsity (but not dimensionality) improvements in harmonic reconstructions.

superiority), we therefore use only the MW sampling theorem for the high-resolution inpainting simulation performed here.

A high-resolution test image is constructed from the same Earth topography data described in Section V-A. Since the original data are defined in harmonic space, we first simply truncate the harmonic coefficients to yield an image band-limited at $L = 128$. In practice, acquired images may not necessarily be band-limited. We thus process the data to construct a test image that is not band-limited. We construct such a test image defined by height above sea-level (*i.e.* we threshold all oceans and trenches, while the continents and mountains remain unaltered). The abrupt transition between the oceans and the continents results in an image that is indeed not band-limited (as verified numerically). Furthermore, the continents and mountainous regions result in a test image that is not highly sparse in its gradient. The resulting realistic test image is shown in Figure 5 (a).

The same measurement procedure as outlined previously is applied to take noisy, incomplete measurements of the data for a range of measurement ratios M/L^2 . The inpainted images are recovered by solving the inpainting problem in harmonic space through (5) using the MW sampling theorem. To solve the inpainting problem for these high-resolution simulations we use the estimator of the inverse transform norm $\|\mathbf{A}\|_2$ described in Section IV-B and the fast adjoint harmonic transform algorithms defined in Section IV-C. Using these fast algorithms, combined with recent optimisations of the SSHT package, it takes approximately 10 minutes to solve the inpainting problem in harmonic space at $L = 128$ on a standard laptop (with a 1.8 GHz Intel Core i7 processor and 4 GB of RAM).

The inpainted images are shown in Figure 5. Since the original realistic test image is not band-limited, the previous SNR measure (which is defined in harmonic space to avoid a dependence on the number of samples of each sampling theorem) is not a meaningful error metric (computation of




Fig. 3. Inpainted images on the sphere recovered by solving the TV inpainting problems for a range of measurement ratios M/L^2 . The ground truth image is shown in Figure 2. The first and second columns of panels show the inpainted images recovered using the DH sampling theorem, while the third and fourth columns show the inpainted images recovered using the MW sampling theorem. The first and third columns of panels show inpainted images recovered by solving the inpainting problem in the spatial domain, while the second and fourth columns show images recovered by solving the inpainting problem in the harmonic domain. The final row of panels corresponds to measurement ratio $M/L^2 = N_{\text{MW}}/L^2 \sim 2$. The quality enhancements due to the MW sampling theorem and by solving the inpainting problem in harmonic space are both clear.

the harmonic transform would indeed be affected by uncontrolled aliasing). Instead, we use the analogous SNR measure defined in image space on the sphere, given by $\text{SNR}_I = 10 \log(\mathbf{x}^\dagger Q \mathbf{x} / ((\mathbf{x}^* - \mathbf{x})^\dagger Q (\mathbf{x}^* - \mathbf{x})))$, where we recall Q is the matrix with quadrature weights on its diagonal. Just as for the definition of the TV norm, the inclusion of the weights for signals that are not band-limited provides only an intuitive approximation to continuous integration. The SNR_I values for each inpainted image are displayed in Figure 5, from which it is apparent that SNR_I increases with increasing measurement ratio. Moreover, it is clearly apparent by eye that our TV inpainting framework is effective when applied to realistic images that are not highly sparse in their gradient and that are not band-limited.

VI. CONCLUSIONS

The MW sampling theorem, developed only recently, achieves a more efficient sampling of the sphere than the standard DH sampling theorem: without any loss to the information content of the sampled signal, the MW sampling theorem reduces the number of samples required to represent a band-limited signal by a factor of two for an equiangular sampling. For signals sparse in a spatially localised measure, such as in a wavelet basis, overcomplete dictionary, or in the magnitude of their gradient, for example, a more efficient sampling enhances the fidelity of sparse image reconstruction through both dimensionality and sparsity. When a signal is recovered directly in the spatial domain, the MW sampling theorem provides enhancements in both dimensionality and sparsity when compared to the DH sampling theorem. By recovering the signal directly in harmonic space it is possible




Fig. 5. Inpainting illustration for a realistic image at high-resolution ($L = 128$). The inpainted images are recovered by solving the inpainting problem in harmonic space using the MW sampling theorem for a range of measurement ratios M/L^2 . The SNR_I of each recovered image is also displayed.

to optimise its dimensionality, in which case the MW sampling theorem still provides an enhancement in sparsity but not in dimensionality.

We verified these statements through a simple inpainting problem on the sphere, where we considered images sparse in their gradient. We built a framework and fast methods for total variation (TV) inpainting on the sphere. Using this framework we performed numerical experiments which confirmed our predictions: in all cases, the more efficient sampling provided by the MW sampling theorem improved the fidelity of sparse image reconstruction on the sphere.

REFERENCES

- [1] N. Jarosik, C. L. Bennett, J. Dunkley, B. Gold, M. R. Greason, M. Halpern, R. S. Hill, G. Hinshaw, A. Kogut, E. Komatsu, D. Larson, M. Limon, S. S. Meyer, M. R. Nolta, N. Odegard, L. Page, K. M. Smith, D. N. Spergel, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright, "Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Sky Maps, Systematic Errors, and Basic Results," *Astrophys. J. Supp.*, vol. 192, p. 14, Feb. 2011.
- [2] H. Johansen-Berg and T. E. J. Behrens, *Diffusion MRI: From quantitative measurement to in-vivo neuroanatomy*. San Diego: Academic Press, 2009.
- [3] R. Ramamoorthi and P. Hanrahan, "A signal processing framework for reflection," *ACM Transactions on Graphics*, vol. 23, no. 4, pp. 1004–1042, 2004.
- [4] E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inform. Theory*, vol. 52, no. 2, pp. 489–509, feb 2006.
- [5] E. Candès, "Compressive sampling," in *Proc. Int. Congress Math.*, ser. Euro. Math. Soc., vol. III, 2006, p. 1433.
- [6] D. Donoho, "Compressed sensing," *IEEE Trans. Inform. Theory*, vol. 52, no. 4, pp. 1289–1306, apr 2006.
- [7] H. Rauhut and R. Ward, "Sparse recovery for spherical harmonic expansions," in *Proc. SampTA*, 2011.
- [8] P. Abrial, Y. Moudden, J.-L. Starck, B. Afeyan, J. Bobin, J. Fadili, and M. K. Nguyen, "Morphological component analysis and inpainting on the sphere: application in physics and astrophysics," *J. Fourier Anal. and Appl.*, vol. 14, no. 6, pp. 729–748, 2007.
- [9] J. D. McEwen and Y. Wiaux, "A novel sampling theorem on the sphere," *IEEE Trans. Sig. Proc.*, vol. 59, no. 12, pp. 5876–5887, 2011.
- [10] J. R. Driscoll and D. M. J. Healy, "Computing Fourier transforms and convolutions on the sphere," *Advances in Applied Mathematics*, vol. 15, pp. 202–250, 1994.
- [11] D. M. J. Healy, D. Rockmore, P. J. Kostelec, and S. S. B. Moore, "FFTs for the 2-sphere – improvements and variations," *J. Fourier Anal. and Appl.*, vol. 9, no. 4, pp. 341–385, 2003.
- [12] E. J. Candès, Y. Eldar, D. Needell, and P. Randall, "Compressed sensing with coherent and redundant dictionaries," *Applied Comput. Harm. Anal.*, vol. 31, no. 1, pp. 59–73, 2010.
- [13] A. Chambolle, "An algorithm for total variation minimization and applications," *J. Math. Imaging Vis.*, vol. 20, no. 1–2, pp. 89–97, 2004.
- [14] P. Combettes and J.-C. Pesquet, "A douglas-rachford splitting approach to nonsmooth convex variational signal recovery," *IEEE J. Selected Top. in Sig. Proc.*, vol. 1, no. 4, pp. 564–574, dec 2007.

- [15] P. Combettes and J.-C. Pesquet, *Proximal splitting methods in signal processing*. New York: Springer, 2011.
- [16] A. Beck and M. Teboulle, "Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems," *IEEE Trans. Image Proc.*, vol. 18, no. 11, pp. 2419–2434, 2009.
- [17] M. J. Fadili and J. L. Starck, "Monotone operator splitting for optimization problems in sparse recovery," in *ICIP*, nov 2009, pp. 1461–1464.
- [18] F. J. Simons, F. A. Dahlen, and M. A. Wiecek, "Spatiospectral concentration on a sphere," *SIAM Rev.*, vol. 48, no. 3, pp. 504–536, 2006.
- [19] D. A. Varshalovich, A. N. Moskalev, and V. K. Khersonskii, *Quantum theory of angular momentum*. Singapore: World Scientific, 1989.

Jason McEwen received a B.E. (Hons) degree in Electrical and Computer Engineering from the University of Canterbury, New Zealand, in 2002 and a Ph.D. degree in Astrophysics from the University of Cambridge in 2006.

He held a Research Fellowship at Clare College, Cambridge, from 2007 to 2008, worked as a Quantitative Analyst from 2008 to 2010, and held a position as a Postdoctoral Researcher at Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland, from 2010 to 2011. From 2011 to 2012 he held a Leverhulme Trust Early Career Fellowship at University College London (UCL), where he remains as a Newton International Fellow, supported by the Royal Society and the British Academy. His research interests are focused on spherical signal processing, including sampling theorems and wavelets on the sphere, compressed sensing and Bayesian statistics, and applications of these theories to cosmology and radio interferometry.

Gilles Puy was born in Nevers, France, on August 8, 1985. He received the Engineering degree from the Ecole Supérieure d'Electricité (Supélec), Gif-sur-Yvette, France, and the M. Sc. degree in Electrical and Electronics Engineering from the Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland, in 2009.

In 2009, he joined the Signal Processing Laboratory and the Laboratory of functional and metabolic imaging at the EPFL as a Doctoral Assistant, where he is currently working towards the Ph.D. degree in Electrical Engineering. His main research interests include inverse problems, compressed sensing, biomedical imaging and radio interferometry.

Jean-Philippe Thiran received the Electrical Engineering and Ph.D. degrees from the Université catholique de Louvain (UCL), Louvain-la-Neuve, Belgium, in 1993 and 1997, respectively.

Since January 2004, he has been an Assistant Professor, responsible for the Image Analysis Group at the Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland. His current scientific interests include image segmentation, prior knowledge integration in image analysis, partial differential equations and variational methods in image analysis, multimodal signal processing, medical image analysis, including multimodal image registration, segmentation, computer-assisted surgery, and diffusion MRI.

Dr. Thiran was Co-Editor-in-Chief of Signal Processing (published by Elsevier Science) from 2001 to 2005. He is currently an Associate Editor of the International Journal of Image and Video Processing (published by Hindawi), and member of the Editorial Board of Signal, Image and Video Processing (published by Springer). He was the General Chairman of the 2008 European Signal Processing Conference (EUSIPCO 2008). He is a senior member of the IEEE, and a member of the MLSP and IVMS technical committees of the IEEE Signal Processing Society.

Pierre Vandergheynst received the M.S. degree in Physics and the Ph.D. degree in Mathematical Physics from the Université catholique de Louvain, Louvain-la-Neuve, Belgium, in 1995 and 1998, respectively.

From 1998 to 2001, he was a Postdoctoral Researcher and an Assistant Professor with the Signal Processing Laboratory, Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland. He is now an Associate Professor at EPFL, where his research focuses on harmonic analysis, sparse approximations, and mathematical image processing with applications to higher dimensional, complex data processing.

Dr. Vandergheynst was co-Editor-in-Chief of Signal Processing from 2002 to 2006 and has been an Associate Editor of the IEEE TRANSACTIONS ON SIGNAL PROCESSING since 2007. He has been on the Technical Committee of various conferences and was Co-General Chairman of the EUSIPCO 2008 conference. He is the author or co-author of more than 50 journal papers, one monograph and several book chapters. He is a laureate of the Apple ARTS award and holds seven patents.

Dimitri Van De Ville received the M.S. degree in Engineering and Computer Sciences from Ghent University, Belgium, in 1998, as well as the Ph.D. degree in Computer Science Engineering, in 2002.

He obtained a grant as Research Assistant with the Fund for Scientific Research Flanders Belgium (FWO). In 2002, he joined Prof. M. Unser's Biomedical Imaging Group at the Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland. In December 2005, he became responsible for the Signal Processing Unit at the University Hospital of Geneva, Geneva, Switzerland, as part of the Centre d'Imagerie Biomédicale (CIBM). He was recently awarded a SNSF professorship from the Swiss National Science Foundation and he currently holds a joint position at the University of Geneva and the EPFL. His research interests include wavelets, sparsity, pattern recognition, and their applications in biomedical imaging, such as functional magnetic resonance imaging.

Dr. Van De Ville served as an Associate Editor for the IEEE Transactions on Image Processing (2006-2009) and the IEEE Signal Processing Letters (2004-2006). He is a member of the Bio Imaging and Signal Processing (BISP) TC of the IEEE SPS. Since 2003, he has also been an Editor and Webmaster of The Wavelet Digest. He is co-chair of the Wavelets series conferences (2007, 2009), together with V. Goyal and M. Papadakis.

Yves Wiaux received the M.S. degree in Physics and the Ph.D. degree in Theoretical Physics from the Université catholique de Louvain (UCL), Louvain-la-Neuve, Belgium, in 1999 and 2002, respectively.

He was a Postdoctoral Researcher at the Signal Processing Laboratories of the Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland, from 2003 to 2008. He was also a Postdoctoral Researcher of the Belgian National Science Foundation (F.R.S.-FNRS) at the Physics Department of UCL from 2005 to 2009. He is now a Maître Assistant of the University of Geneva (UniGE), Switzerland, with joint affiliation between the Institute of Electrical Engineering and the Institute of Bioengineering of EPFL, and the Department of Radiology and Medical Informatics of UniGE. His research lies at the intersection between complex data processing (including development on wavelets and compressed sensing) and applications in astrophysics (notably in cosmology and radio astronomy) and in biomedical sciences (notably in MRI and diffusion MRI).